Mean-Semivariance Optimization: A Heuristic Approach

Javier Estrada

Academics and practitioners optimize portfolios using the mean-variance approach far more often than the mean-semivariance approach, despite the fact that semivariance is often considered a more plausible measure of risk than variance. The popularity of the mean-variance approach follows in part from the fact that mean-variance problems have well-known closed-form solutions, whereas mean-semivariance optimal portfolios cannot be determined without resorting to obscure numerical algorithms. This follows from the fact that, unlike the exogenous covariance matrix, the semicovariance matrix is endogenous. This article proposes a heuristic approach that yields a symmetric and exogenous semicovariance matrix, which enables the determination of mean-semivariance optimal portfolios by using the well-known closed-form solutions of mean-variance problems. The heuristic proposed is shown to be both simple and accurate.

■As is well known, Markowitz (1952) pioneered the issue of portfolio optimization with a seminal article, later expanded into a seminal book (Markowitz, 1959). Also well known is that at the heart of the portfolio-optimization problem, there is an investor whose utility depends on the expected return and risk of his portfolio, the latter quantified by the variance of returns.

What may be less well known is that, from the very beginning, Markowitz favored another measure of risk: the semivariance of returns. In fact, Markowitz (1959) allocates the entire chapter IX to discuss semivariance, where he argues that "analyses based on *S* [semivariance] tend to produce better portfolios than those based on *V* [variance]" (see Markowitz, 1991, page 194). In the revised edition of his book (Markowitz, 1991), he goes further and claims that "semivariance is the more plausible measure of risk" (page 374). Later he claims that because "an investor worries about underperformance rather than overperformance, semideviation is a more appropriate measure of investor's risk than variance" (Markowitz, Todd, Xu, and Yamane, 1993, page 307).

Why, then, have practitioners and academics been optimizing portfolios for more than 50 years using variance as a measure of risk? Simply because, as Markowitz (1959) himself suggests, variance has an edge over semivariance "with respect to cost, convenience, and familiarity" (see Markowitz, 1991, page 193). He therefore focused his analysis on variance, practitioners, and academics followed his lead, and the rest is history.

Familiarity, however, has become less of an issue over time. In fact, in both practice and academia, downside risk has been gaining increasing attention, and the many magnitudes that

---

Javier Estrada is a Professor of Finance at the IESE Business School in Barcelona, Spain.

I would like to thank David Nawrocki and Charles Smithson for their comments. Gabriela Giannattasio provided valuable research assistance. The views expressed below and any errors that may remain are entirely my own. The heuristic proposed in this article has been incorporated into the downside risk Excel add-in provided by "hoadley Trading & Investment Tools."

capture downside risk are by now well known and widely used. The focus of this article is on the issues of cost and convenience.

The difference in cost, Markowitz (1959) argues, is given by the fact that efficient sets based on semivariance took, back then, two to four times as much computing time as those based on variance. The difference in convenience, in turn, is given by the fact that efficient sets based on variance require as inputs only means, variances, and covariances, whereas those based on semivariance require the entire joint distribution of returns. The ultimate goal of this article, then, is to propose a heuristic approach to the estimation of portfolio semivariance that renders the issues of cost and convenience irrelevant, thus hopefully removing the last remaining obstacles to a widespread use of mean-semivariance optimization.

In a nutshell, this article proposes to estimate the semivariance of portfolio returns by using an expression similar to that used to estimate the variance of portfolio returns. The advantages of this approach are twofold: 1) estimating the semivariance of portfolio returns is just as easy as estimating the variance of portfolio returns (and in both cases the same number of inputs is required); and 2) it can be done with an expression well known by all practitioners and academics, without having to resort to any black-box numerical algorithm. In addition, the heuristic proposed here yields a portfolio semivariance that is both very highly correlated and very close in value to the exact magnitude it intends to approximate.

The article is organized as follows. Section I introduces the issue, discusses the difficulties related to the optimization of portfolios on the basis of means and semivariances, and shows how they are overcome by the heuristic approach proposed in this article. Section II, based on data on individual stocks, markets, and asset classes, provides empirical support for this heuristic. Section III concludes with an assessment.

I. The Issue

There is little doubt that practitioners rely much more on mean-variance optimization than on mean-semivariance optimization. This is largely because, unlike the neat closed-form solutions of mean-variance problems known by most academics and practitioners, mean-semivariance problems are usually solved with obscure numerical algorithms. This, in turn, is largely because, unlike the exogenous covariance matrix used in mean-variance problems, the semicovariance matrix of mean-semivariance problems is, as will be illustrated, endogenous.

This section starts with some basic definitions and notation, and then introduces the definition of portfolio semivariance proposed in this article. A numerical example is then used to illustrate both the endogeneity of the usual definition of the semicovariance matrix and the exogeneity of the definition proposed here.

A. The Basics

Consider an asset i with returns R_{it}, where t indexes time. The variance of this asset's returns (σ_i^2) is given by

$$\sigma_i^2 = \mathrm{E}[(R_i - \mu_i)^2] = (1/T)\cdot\sum_{t=1}^{T}(R_{it} - \mu_i)^2 \ , \qquad (1)$$

where μ_i denotes the mean return of asset i and T the number of observations; and the covariance between two assets i and j (σ_{ij}) is given by

$$\sigma_{ij} = \mathrm{E}[(R_i - \mu_i)(R_j - \mu_j)] = (1/T)\cdot\sum_{t=1}^{T}(R_{it} - \mu_i)(R_{jt} - \mu_j)\ . \qquad (2)$$

The semivariance of asset i's returns with respect to a benchmark B (Σ_{iB}^2) is given by

$$\Sigma_{iB}^2 = \mathrm{E}\{[\mathrm{Min}(R_i - B, 0)]^2\} = (1/T)\cdot\sum_{t=1}^{T}[\mathrm{Min}(R_{it} - B, 0)]^2\ , \qquad (3)$$

where B is any benchmark return chosen by the investor. The square root of Equation (3) is the semideviation of asset i with respect to benchmark B, a widely-used measure of downside risk. Section A of the Appendix,[1] borrowing heavily from Estrada (2006), provides a very brief introduction to the semideviation and discusses some of the advantages it has over the standard deviation as a measure of risk; Nawrocki (1999) provides a brief history of downside risk and an overview of downside risk measures.

The semicovariance between assets i and j (Σ_{ij}) is trickier to define. Hogan and Warren (1974) define it as

$$\Sigma_{ij}^{HW} = \mathrm{E}\{(R_i - R_f)\cdot \mathrm{Min}(R_j - R_f, 0)\}\ , \qquad (4)$$

where the superscript HW indicates that this is definition proposed by Hogan and Warren. This definition, however, has two drawbacks: 1) the benchmark return is limited to the risk-free rate and cannot be tailored to any desired benchmark, and 2) it is usually the case that $\Sigma_{ij}^{HW} \neq \Sigma_{ji}^{HW}$. This second characteristic is particularly limiting both formally (the semicovariance matrix is usually asymmetric) and intuitively (it is not clear how to interpret the contribution of assets i and j to the risk of a portfolio).

In order to overcome these two drawbacks, Estrada (2002, 2007) defines the semicovariance between assets i and j with respect to a benchmark B (Σ_{ijB}) as

$$\begin{aligned}\Sigma_{ijB} &= \mathrm{E}\{\mathrm{Min}(R_i - B, 0)\cdot \mathrm{Min}(R_j - B, 0)\} = \\ &(1/T)\cdot\sum_{t=1}^{T}[\mathrm{Min}(R_{it} - B, 0)\cdot \mathrm{Min}(R_{jt} - B, 0)]\end{aligned}\ . \qquad (5)$$

This definition can be tailored to any desired B and generates a symmetric ($\Sigma_{ijB}=\Sigma_{jiB}$) and, as will be shown, exogenous semicovariance matrix. Both the symmetry and exogeneity of this matrix are critical for the implementation of the proposed heuristic.

Finally, the expected return (E_p) and variance (σ_p^2) of a portfolio are given by

$$E_p = \sum_{i=1}^{n} x_i E_i \ , \quad (6)$$

$$\sigma_p^2 = \sum_{i=1}^{n} \sum_{j=1}^{n} x_i x_j \sigma_{ij} \ , \quad (7)$$

where x_i denotes the proportion of the portfolio invested in asset i, E_i the expected return of asset i, and n the number of assets in the portfolio.

B. The Problem

Portfolio-optimization problems can be specified in many ways depending on the goal and restrictions of the investor.[1] The problem of minimizing the risk of a portfolio subject to a target return (E^T) is given by

$$\text{Min}_{x_1, x_2, \ldots, x_n} \quad \sigma_p^2 = \sum_{i=1}^{n} \sum_{j=1}^{n} x_i x_j \sigma_{ij} \quad (8)$$

$$\sum_{i=1}^{n} x_i E_i = E^T \quad \text{and} \quad \sum_{i=1}^{n} x_i = 1 \ , \quad (9)$$

where risk is measured as the variance of portfolio returns.

This problem can be solved for a specific value of E^T or, alternatively, for several values of E^T thus generating the minimum-variance set. Either way, it is important to notice, first, that the risk of the portfolio can be expressed as a function of the risk of the individual assets in the portfolio; second, that all the variances, covariances, and expected returns of the individual assets are exogenous variables; and third, that this problem has a well-known closed-form solution. For this reason, although it is not important for the purposes of this article how the σ_{ij} are estimated, it *is* important that, once the values of these parameters (exogenous variables) are determined, they become inputs (together with E_i and E^T) in the closed-form solution of the problem, which in turn yields the optimal allocation to each of the n assets in the portfolio (the endogenous variables x_i).

But what if, instead of defining risk as the variance of portfolio returns, an investor wanted to define it as the *semivariance* of portfolio returns? What if, given a benchmark return B chosen by the investor, he wanted to

$$\text{Min}_{x_1, x_2, \ldots, x_n} \quad \Sigma_{pB}^2 = (1/T) \cdot \sum_{t=1}^{T} [\text{Min}(R_{pt} - B, 0)]^2 \quad (10)$$

$$\sum_{i=1}^{n} x_i E_i = E^T, \quad \sum_{i=1}^{n} x_i = 1, \quad \text{and} \quad x_i \geq 0 \ , \quad (11)$$

where R_{pt} denotes the returns of the portfolio and Σ_{pB}^2 their semivariance? The main obstacle to the solution of this problem is that the semicovariance matrix is endogenous; that is, a change in weights affects the periods in which the portfolio underperforms the benchmark, which in turn affects the elements of the semicovariance matrix.[2]

In order to overcome this obstacle, many algorithms have been proposed to solve the problem in Equations (10) and (11), some of which are discussed below. More importantly, this article proposes a heuristic approach to solve this problem without having to resort to any black-box numerical algorithm. In fact, as will be evident, the heuristic proposed here makes it possible to solve not only the problem in Equations (10) and (11), but also all mean-semivariance problems with the same well-known closed-form solutions widely-used to solve mean-variance problems.

More precisely, this article argues that the semivariance of a portfolio with respect to a benchmark B can be approximated with the expression

$$\Sigma_{pB}^2 \approx \sum_{i=1}^{n} \sum_{j=1}^{n} x_i x_j \Sigma_{ijB} \ , \quad (12)$$

where Σ_{ijB} is defined as in Equation (5). This expression yields a symmetric and exogenous semicovariance matrix, which can then be used in the same way the (symmetric and exogenous) covariance matrix is used in the solution of mean-variance problems.

C. An Example

Table I displays the annual returns of the S&P-500 and the Nikkei-225 between 1997 and 2006, as well as the return of two portfolios: one invested 80% in the S&P and 20% in the Nikkei, and the other invested 10% in the S&P and 90% in the Nikkei. Consider for now the 80-20 portfolio. The standard deviation of this portfolio can be calculated by first estimating its returns over the sample period, and then calculating the standard deviation of those returns. The fourth column of Table I shows the returns of the 80-20 portfolio, and the standard deviation of those returns, which can be straightforwardly calculated using the square root of (1), is 16.7%.

Importantly, the standard deviation of the 80-20 portfolio can also be calculated by using the square root of Equation (7). Taking into account that the standard deviations of the S&P and the Nikkei over the 1997-2006 period are 17.8%

[1] There are four standard portfolio optimization problems: 1) minimizing the risk of a portfolio; 2) minimizing the risk of a portfolio subject to a target return; 3) maximizing the return of a portfolio subject to a target level of risk; and 4) maximizing the risk-adjusted return of a portfolio. The heuristic proposed here applies to all four problems; only, for concreteness, most of the discussion is focused on the second problem.

[2] Note, also, that this formulation of the optimization problem ignores both the downside risk of individual assets and the downside covariance between individual assets; see Sing and Ong (2000).

Table I. The Endogenous Semicovariance Matrix

This table shows the returns over the 1997-2006 period of the S&P-500 (S&P), the Nikkei-225 (Nikkei), a portfolio invested 80% in the S&P and 20% in the Nikkei (80-20), and a portfolio invested 10% in the S&P and 90% in the Nikkei (10-90). 'Conditional returns' are defined as 0% when the return of the portfolio is positive, and the return of the asset when the return of the portfolio is negative. All returns are in dollars and account for capital gains and dividends. All numbers are in percentages.

Year	S&P	Nikkei	80-20	10-90	'Conditional Returns'					
					80-20 Portfolio			10-90 Portfolio		
					S&P	Nikkei	Product	S&P	Nikkei	Product
1997	31.0	–21.2	20.6	–16.0	0.0	0.0	0.0	31.0	–21.2	–6.6
1998	26.7	– 9.3	19.5	–5.7	0.0	0.0	0.0	26.7	–9.3	–2.5
1999	19.5	36.8	23.0	35.1	0.0	0.0	0.0	0.0	0.0	0.0
2000	–10.1	–27.2	–13.5	–25.5	–10.1	–27.2	2.8	–10.1	–27.2	2.8
2001	–13.0	–23.5	–15.1	–22.5	–13.0	–23.5	3.1	–13.0	–23.5	3.1
2002	–23.4	–18.6	–22.4	–19.1	–23.4	–18.6	4.4	–23.4	–18.6	4.4
2003	26.4	24.5	26.0	24.6	0.0	0.0	0.0	0.0	0.0	0.0
2004	9.0	7.6	8.7	7.7	0.0	0.0	0.0	0.0	0.0	0.0
2005	3.0	40.2	10.4	36.5	0.0	0.0	0.0	0.0	0.0	0.0
2006	13.6	6.9	12.3	7.6	0.0	0.0	0.0	0.0	0.0	0.0

and 24.1%, and that the covariance between these two indices is 0.0163, it follows from Equation (7) that

$$\sigma_p = \{(0.82)(0.1782) + (0.22)(0.2412) + 2(0.8)(0.2)(0.0163)\}1/2 = 16.7\% ,$$

which is, of course, identical to the number obtained before from the portfolio returns. So far, no mystery here.

The problem arises if the proper measure of risk is not the portfolio's variance, but its semivariance. One obvious way of calculating this magnitude would be by first calculating the returns of the portfolio and then using Equation (3) to calculate the semivariance of its returns. Assume a benchmark return of 0% (B=0), and consider again the 80-20 portfolio. We could first calculate the returns of this portfolio (shown in the fourth column of Table I), and then calculate the semivariance of its returns by using Equation (3). That would obtain a portfolio semivariance with respect to 0% equal to 0.0092, and a portfolio semideviation equal to $(0.0092)^{1/2}$=9.6%.

Thus, for any given portfolio, its semideviation can always be calculated as just explained. But here is the problem: if instead of the semideviation of one portfolio, we wanted to calculate the portfolio with the lowest semideviation from a set of, say, 1,000 feasible portfolios, we would first need to calculate the returns of each portfolio; then from those returns we would need to calculate the semideviation of each portfolio; and finally from those semideviations we would need to select the one with the lowest value. Obviously, as the number of assets in the portfolio increases, and the number of feasible portfolios increases even more, choosing the optimal portfolio with this procedure becomes intractable.

Look at this from a different perspective. If the elements of the semicovariance matrix were exogenous, then we could formally solve the given optimization problem and obtain a closed-form solution. We could then input into this closed-form solution the values of the exogenous variables of the problem at hand, and obtain as a result the weights that satisfy the problem. This is exactly what investors routinely do when solving portfolio-optimization problems in the mean-variance world. But the problem in the mean-semivariance world is, precisely, that the elements of the semicovariance matrix are not exogenous.

D. The Endogeneity of the Semicovariance Matrix

Markowitz (1959) suggests estimating the semivariance of a portfolio with the expression

$$\Sigma_{pB}^2 = \sum_{i=1}^{n} \sum_{j=1}^{n} x_i x_j S_{ijB} \ , \tag{13}$$

where

$$S_{ijB} = (1/T) \cdot \sum_{t=1}^{K} (R_{it} - B)(R_{jt} - B) \ , \tag{14}$$

where periods 1 through K are those in which *the portfolio* underperforms the benchmark return B.

This definition of portfolio semivariance has one advantage and one drawback. The advantage is that it provides an exact estimation of the portfolio semivariance. The drawback is that the semicovariance matrix is endogenous; that is, a change in weights affects the periods in which the portfolio underperforms the benchmark, which in turn affects the elements of the semicovariance matrix.

To see the advantage of this definition of portfolio semivariance, go back to the 80-20 portfolio in Table I, and consider again B=0. The sixth column of this table shows the 'conditional returns' of the S&P defined, following Equation (14), as 0% when the return of the 80-20 portfolio is positive (thus outperforming the benchmark), and the return of the S&P when the return of the 80-20 portfolio is negative (thus underperforming the benchmark). To illustrate, the 'conditional return' of the S&P is 0% in 1997 because the 80-20 portfolio delivered a positive return, and –10.1% (the return of the S&P) in 2000 because the 80-20 portfolio delivered a negative return. The seventh column shows the 'conditional returns' for the Nikkei, and the eighth column is just the product of the sixth and the seventh columns.

The heuristic proposed is both simple and accurate. Estimating semicovariances is just as easy as estimating covariances, and aggregating them into a portfolio semivariance is, with the proposed heuristic, just as easy as aggregating covariances into a portfolio variance.

The four terms of the semicovariance matrix that follow from Equation (14) can be calculated as follows. Squaring the 'conditional returns' in the sixth column and taking their average obtains $S_{S\&P,S\&P,0}$=0.0082; doing the same with the 'conditional returns' in the seventh column obtains $S_{Nikkei,Nikkei,0}$=0.0164; and taking the average of the numbers in the eighth column obtains $S_{S\&P,Nikkei,0}$=0.0102. Then, it follows from Equation (13) that the semivariance of the 80-20 portfolio is

$$\{(0.8^2)(0.0082) + (0.2^2)(0.0164) + 2(0.8)(0.2)(0.0102)\} = 0.0092,$$

and its semideviation is $(0.0092)^{1/2}$=9.6%, which is exactly the same number obtained before.

Therefore, the expression proposed by Markowitz (1959) does indeed provide an exact estimation of the portfolio semivariance. But the problem is that, in order to estimate this semivariance, we need to know whether the *portfolio* underperforms the benchmark, and we then run into the problem previously mentioned: the semicovariance matrix is endogenous because a change in weights affects when the portfolio underperforms the benchmark, which in turn affects the elements of the semicovariance matrix.

To see this more clearly, go back to Table I and consider now the portfolio invested 10% in the S&P and 90% in the Nikkei. The returns of this portfolio are shown in the fifth column, the 'conditional returns' (as defined above) of the S&P and the Nikkei in the ninth and tenth columns, and the product of these last two columns in the eleventh column. Importantly, note that the 'conditional returns' of the S&P and the Nikkei for the 10-90 portfolio that follow from Equation (14) are different from those for the 80-20 portfolio that follow from the same expression.

The four terms of the semicovariance matrix that follows from Equation (14) can be calculated as before. Squaring the numbers in the ninth column and then taking their average obtains $S_{S\&P,S\&P,0}$=0.0249; squaring the numbers in the tenth column and then taking their average obtains $S_{Nikkei,Nikkei,0}$=0.0217; and taking the average of the numbers in the last column obtains $S_{S\&P,Nikkei,0}$=0.0011. And importantly, *note that all these numbers are different from those calculated for the 80-20 portfolio*. This clearly illustrates that the semicovariance matrix is endogenous because its elements depend on the asset weights.

Finally, for the sake of completeness, with the numbers just calculated Equation (13) can be used to calculate the semivariance of the 10-90 portfolio, which is given by

$$\{(0.1^2)(0.0249) + (0.9^2)(0.0217) + 2(0.1)(0.9)(0.0011)\} = 0.0181,$$

thus implying a semideviation of $(0.0181)^{1/2}$=13.4%.

E. Some Possible Solutions

The endogeneity of the semicovariance matrix as defined in Equation (14) has led many authors to propose different ways of tackling the problem in Equations (10) and (11). Hogan and Warren (1972) propose to solve this problem using the Frank-Wolfe algorithm; they explain the two basic steps of this iterative method (the direction-finding problem and the step-size problem) and illustrate its application with a simple hypothetical example. Ang (1975) proposes to linearize the semivariance so that the optimization problem can be solved using linear (instead of quadratic) programming.

Nawrocki (1983) proposes a further simplification of the heuristic proposed by Elton, Gruber, and Padberg (1976). The latter focus on the mean-variance problem and impose the simplifying assumption that all pairwise correlations are the same; the former further imposes a value of zero for all of these correlations and extends the analysis to other measures of risk, including the semivariance. In this heuristic, assets are ranked according to the measure $z_i = (E_i – R_f)/RM_i$, where RM_i is a risk measure for asset i, and assets with z_i>0 are included in the portfolio according to the proportions $w_i = z_i/\Sigma_i z_i$. Nawrocki and Staples (1989) expand the scope of Nawrocki (1983) by considering the lower partial moment (LPM) as a risk measure.

Harlow (1991) also considers the problem in Equations (10) and (11) and generates mean-semivariance efficient frontiers, which he compares to mean-variance efficient frontiers. However, he does not explain how these frontiers are obtained other than stating that the optimization process uses the entire distribution of returns. Similarly, Grootveld and Hallerbach (1999) generate mean-LPM efficient frontiers and state that the numerical optimization process they use for solving the problem in Equations (10) and (11) is tedious and demanding, but do not provide details of such process.

Markowitz et al. (1993) transform the mean-semivariance problem into a quadratic problem by adding fictitious securities. This modification enables them to apply to the modified mean-semivariance problem the critical line algorithm originally developed to solve the mean-variance problem.

More recently, de Athayde (2001) proposes a non-parametric approach to calculate the portfolio semivariance, as well as an algorithm (basically a series of standard minimization problems) to optimize it and generate the efficient frontier. Ballestero (2005), in turn, proposes a definition of portfolio semivariance (restricting the benchmark to the mean) that, when incorporated into optimization problems, these can be solved by applying parametric quadratic programming methods.

F. A Heuristic Approach

As previously advanced, the heuristic proposed in this article is based on estimating the portfolio semivariance using Equation (12), which in turn is based on Equation (5), which generates a symmetric and exogenous semicovariance matrix. Recall that with Equation (14) knowledge of whether *the portfolio* underperforms the benchmark B is needed, which generates the endogeneity problem discussed earlier. With Equation (5), however, knowledge of whether *the asset* (not the portfolio) underperforms the benchmark B is needed. Again, an example may help.

Table II reproduces the returns over the 1997-2006 period of the S&P, the Nikkei, the 80-20 portfolio, and the 10-90 portfolio, all taken from Table I. As previously illustrated, the elements of the semicovariance matrix that follow from Equation (14) for the 80-20 portfolio are different from those of the semicovariance matrix that follow from Equation (14) for the 10-90 portfolio, which confirms the endogeneity of this definition of semicovariance. As will be shown, the elements of the semicovariance matrix that follow from Equation (5) are invariant to the portfolio considered and are, therefore, exogenous.

To see this, calculate the four terms of the semicovariance matrix that follow from this expression by considering once again a benchmark return of 0%. First, redefine 'conditional returns' as 0% when the return of *the asset* is positive (thus outperforming the benchmark), and the return of the asset when the return of *the asset* is negative (thus underperforming the benchmark). To illustrate, the conditional return of the S&P is 0% in 1997 because *the S&P* delivered a positive return, and –10.1% (the return of the S&P) in 2000 because *the S&P* delivered a negative return.

These 'conditional returns' of the S&P and the Nikkei are shown in the sixth and seventh columns of Table II, and the eighth column is the product of the previous two. Note that because these 'conditional returns' depend on whether *the asset*, not the portfolio, underperforms the benchmark, they are relevant not only to estimate the semicovariance matrix of the 80-20 portfolio, but also that of any other portfolio.

The four terms of the semicovariance matrix that follow from Equation (5), then, can be calculated as follows. Squaring the 'conditional returns' in the sixth column and taking their average obtains $\Sigma_{S\&P,S\&P,0}$=0.0082; doing the same with the 'conditional returns' of the seventh column obtains $\Sigma_{Nikkei,Nikkei,0}$=0.0217; and taking the average of the numbers in the eighth column obtains $\Sigma_{S\&P,Nikkei,0}$=0.0102. Then, it follows from Equation (12) that the semivariance of the 80-20 portfolio is

$$\{(0.8^2)(0.0082) + (0.2^2)(0.0217) + 2(0.8)(0.2)(0.0102)\} = 0.0094,$$

and its semideviation is $(0.0094)^{1/2}$=9.7%, very close to the exact 9.6% number calculated previously from the portfolio returns.

Importantly, if Equation (12) is used to calculate the semivariance of the 10-90 portfolio, then

$$\{(0.1^2)(0.0082) + (0.9^2)(0.0217) + 2(0.1)(0.9)(0.0102)\} = 0.0195,$$

thus implying a semideviation of $(0.0195)^{1/2}$=14.0%. Note that this number is very close to the exact 13.4% figure calculated for this portfolio in section I.D. More importantly, note that the only difference between this calculation and that for the 80-20 portfolio is in the weights; *the four elements of the semicovariance matrix are the same*.

In short, if semicovariances are defined as in Equation (14) and portfolio semivariance as in Equation (13), then the endogeneity problem occurs and black-box numerical algorithms need to be used to solve portfolio-optimization problems. If semicovariances are instead defined as in Equation (5) and portfolio semivariance as in Equation (12), then a symmetric and exogenous semicovariance matrix is obtained, and the well-known and widely-used closed-form

Table II. The Exogenous Semicovariance Matrix

This table shows the returns over the 1997-2006 period of the S&P-500 (S&P), the Nikkei-225 (Nikkei), a portfolio invested 80% in the S&P and 20% in the Nikkei (80-20), and a portfolio invested 10% in the S&P and 90% in the Nikkei (10-90). 'Conditional returns' are defined as 0% when the return of the asset is positive, and the return of the asset when the return of the asset is negative. All returns are in dollars and account for capital gains and dividends. All numbers are in percentages.

Year	Assets		Portfolios		'Conditional Returns'		
	S&P	Nikkei	80-20	10-90	S&P	Nikkei	Product
1997	31.0	–21.2	20.6	–16.0	0.0	–21.2	0.0
1998	26.7	–9.3	19.5	–5.7	0.0	–9.3	0.0
1999	19.5	36.8	23.0	35.1	0.0	0.0	0.0
2000	–10.1	–27.2	–13.5	–25.5	–10.1	–27.2	2.8
2001	–13.0	–23.5	–15.1	–22.5	–13.0	–23.5	3.1
2002	–23.4	–18.6	–22.4	–19.1	–23.4	–18.6	4.4
2003	26.4	24.5	26.0	24.6	0.0	0.0	0.0
2004	9.0	7.6	8.7	7.7	0.0	0.0	0.0
2005	3.0	40.2	10.4	36.5	0.0	0.0	0.0
2006	13.6	6.9	12.3	7.6	0.0	0.0	0.0

solutions of mean-variance portfolio optimization problems can be applied.

G. A First Look at the Accuracy of the Approximation

In order to take a preliminary look at the accuracy of the approximation proposed, and to round up the example discussed so far, Table III shows the returns of eleven portfolios over the 1997-2006 period that differ only in the proportions invested in the S&P and the Nikkei. The third row from the bottom shows the exact semideviation of each portfolio calculated from the portfolio returns and based on Equation (3), and the second row from the bottom shows the semideviation of each portfolio based on the approximation proposed in Equation (12). In both cases the benchmark return B is 0%. The last row shows the difference between the exact and the approximate semideviations.

The correlation between the exact semideviations based on Equation (3) and the approximate semideviations based on Equation (12) is a whopping 0.98. Furthermore, the difference between the approximate and the exact semideviations is under 1% in all cases, with an average of 0.42%. Finally, the direction of the error is predictable; whenever there is a difference between the two, the approximate semideviation is larger than its exact counterpart. In other words, whenever the approximation errs, it does so on the side of caution, overestimating (by a small amount) the risk of the portfolio.

II. The Evidence

The heuristic proposed in this article yields a symmetric and exogenous semicovariance matrix which, as discussed earlier, makes it possible to solve mean-semivariance optimization problems using the well-known closed-form solutions widely-used for mean-variance optimization problems. However, as with any heuristic, its usefulness rests on its simplicity and accuracy. Its simplicity is hopefully evident from the previous discussion; its accuracy is discussed next.

This section starts by considering portfolios of stocks, markets, and asset classes with the purpose of comparing their exact semideviations to the approximate semideviations based on the proposed heuristic. It then considers mean-variance and mean-semivariance optimal portfolios, the latter based on the proposed heuristic, with the goal of comparing the allocations generated by these two approaches.

A. The Accuracy of the Approximation

In order to test the accuracy of the proposed heuristic over a wide range of assets, exact and approximate semideviations were calculated for over 1,100 portfolios, some containing stocks, some markets, and some asset classes. The data is described in detail in section A of the Appendix. Table IV summarizes the results of all the estimations.

Panel A shows the results for two-asset portfolios selected from three asset classes: 1) US stocks, 2) emerging markets stocks, and 3) US real estate, all of which exhibit statistically-significant negative skewness over the sample period.

Table III. A First Look at the Accuracy of the Approximation

This table shows the return of eleven portfolios over the 1997-2006 period, each with different proportions invested in the S&P-500 (S&P) and the Nikkei-225 (Nikkei). The returns of the S&P and the Nikkei are those in Tables I and II. Σ_{p0} Equation (3) and Σ_{p0} Equation (12) denote the portfolio semideviations based on Equations (3) and (12), both with respect to a benchmark return of 0%. The last row shows the differences Σ_{p0} Equation (12) – Σ_{p0} Equation (3). All returns are in dollars and account for capital gains and dividends. All numbers are in percentages.

Year	Proportion of the Portfolio Invested in the S&P (%)										
	100	90	80	70	60	50	40	30	20	10	0
1997	31.0	25.8	20.6	15.3	10.1	4.9	–0.3	–5.5	–10.8	–16.0	–21.2
1998	26.7	23.1	19.5	15.9	12.3	8.7	5.1	1.5	–2.1	–5.7	–9.3
1999	19.5	21.3	23.0	24.7	26.4	28.2	29.9	31.6	33.3	35.1	36.8
2000	–10.1	–11.8	–13.5	–15.3	–17.0	–18.7	–20.4	–22.1	–23.8	–25.5	–27.2
2001	–13.0	–14.1	–15.1	–16.2	–17.2	–18.3	–19.3	–20.4	–21.4	–22.5	–23.5
2002	–23.4	–22.9	–22.4	–21.9	–21.5	–21.0	–20.5	–20.0	–19.6	–19.1	–18.6
2003	26.4	26.2	26.0	25.8	25.6	25.4	25.2	25.0	24.8	24.6	24.5
2004	9.0	8.9	8.7	8.6	8.4	8.3	8.2	8.0	7.9	7.7	7.6
2005	3.0	6.7	10.4	14.2	17.9	21.6	25.3	29.1	32.8	36.5	40.2
2006	13.6	12.9	12.3	11.6	10.9	10.3	9.6	8.9	8.3	7.6	6.9
Σ_{p0} Equation (3)	9.05	9.29	9.57	9.88	10.23	10.60	11.00	11.56	12.36	13.44	14.75
Σ_{p0} Equation (12)	9.05	9.32	9.68	10.12	10.64	11.21	11.84	12.52	13.23	13.98	14.75
Difference	0.00	0.03	0.11	0.24	0.41	0.61	0.84	0.96	0.87	0.53	0.00

Table IV. Exact and Approximate Portfolio Semideviations

This table shows the annualized semideviations of portfolios with respect to a benchmark return of 0% over the January 1997-December 2006 period. For the assets in each line of Panel A, 101 portfolios were generated, with weights varying between 0% and 100% in each asset (in increments of 1%), and their semideviation calculated. Avg Σ_{p0} Equation (3) and Avg Σ_{p0} Equation (12) denote the average portfolio semideviations across the 101 portfolios based on Equations (3) and (12). Σ_{p0}Equation (3) Range denotes the range between the minimum and the maximum values of Σ_{p0} Equation (3) across the 101 portfolios. Difference is between Avg Σ_{p0} Equation (12) and Avg Σ_{p0} Equation (3) and Rho denotes the correlation betweenΣ_{p0} Equation (3) and Σ_{p0} Equation (12), both across the 101 portfolios. For the assets in each line of Panels B, C, and D, 100 random portfolios were generated and the process outlined for Panel A was repeated. All returns are monthly, in dollars, and account for capital gains and dividends. All numbers but correlations are in percentages. A full data description is available in the appendix.

	Σ_{p0} Equation (3) Range	Avg. Σ_{p0} Equation (3)	Avg. Σ_{p0} Equation (12)	Difference	Rho
Panel A. Asset Classes					
USA-EMI	10.21–17.29	13.15	13.24	0.09	1.00
USA-NAREIT	7.37–10.21	8.16	8.48	0.32	0.99
EMI-NAREIT	8.48–17.29	11.66	11.96	0.30	1.00
Panel B. Emerging Markets					
Group 1 (5 markets)	14.75–19.70	16.80	18.90	2.10	0.97
Group 2 (5 markets)	15.24–17.51	16.06	16.74	0.68	0.98
Group 3 (10 markets)	13.92–16.38	15.11	16.91	1.80	0.93
Panel C. DJIA Stocks					
Group 1 (10 stocks)	9.99–13.16	11.26	13.33	2.07	0.90
Group 2 (10 stocks)	9.12–14.87	12.31	14.76	2.45	0.99
Group 3 (10 stocks)	9.34–12.37	10.59	12.92	2.34	0.95
Group 4 (30 stocks)	9.39–11.16	10.23	12.90	2.67	0.91
Panel D. Asset Classes					
5 Asset classes	4.32–11.66	7.43	8.30	0.87	1.00

Consider first the combination of US stocks and emerging markets stocks (first line of Panel A). Portfolios were formed with weights varying between 0% and 100% in the US market (the rest being allocated to emerging markets), in increments of 1%. Monthly returns over the January 1997-December 2006 period were then calculated for these 101 portfolios. Using these returns, the exact semideviation with respect to a 0% benchmark return was calculated for all these portfolios according to Equation (3) and subsequently annualized. The second column of the table shows the minimum (10.21%) and maximum (17.29%) values across the 101 annualized semideviations, and the third column shows the average (13.15%).

Semivariance is a more plausible measure of risk than variance, as Markowitz (1991) himself suggested, and the heuristic proposed here makes mean-semivariance optimization just as easy to implement as mean-variance optimization.

Approximate semideviations according to Equation (12), with respect to a 0% benchmark return, were then calculated and subsequently annualized for the 101 portfolios. The average among all these annualized approximate semideviations (13.24%) is reported in the fourth column of the table. The difference between the average exact semideviation and the average approximate semideviation is reported in the fifth column, and at 0.09% in annual terms is basically negligible. Furthermore, the correlation between the 101 exact and approximate semideviations, reported in the last column, is a perfect 1.00. These results obviously support the heuristic proposed here.

The results are equally encouraging for portfolios of the other two-asset combinations in Panel A, generated with the methodology already described. The average difference between exact and approximate semideviations is 0.32% for the 101 portfolios of US stocks and real estate, and 0.30% for the 101 portfolios of emerging markets and real estate, both in annual terms. The correlations between the 101 exact and approximate semideviations are 0.99 in the first case and 1.00 in the second. Again, these results are clearly encouraging.

Panel B shows portfolios of emerging markets. Group 1 consists of 5 emerging markets (China, Egypt, Korea, Malaysia, and Venezuela) that over the sample period displayed statistically-significant positive skewness. Portfolios were formed by generating 100 random weights for each of these indices, subsequently standardized to ensure that for each portfolio their sum added to one. As before, returns for these 100 portfolios were calculated over the January 1997-December 2006 period. Then, exact and approximate semideviations with respect to a 0% benchmark return were calculated for all portfolios and subsequently annualized. The correlation between the 100 exact and approximate semideviations is 0.97 and the difference between the averages is in this case higher, 2.1% in annual terms.

Group 2 consists of five emerging markets (Chile, Hungary, Mexico, Peru, and South Africa) that over the sample period displayed statistically-significant negative skewness. Portfolios combining these five markets were generated with the same methodology described for the markets in Group 1. As the table shows, the correlation between the 100 exact and approximate semideviations is 0.98, and the difference between the averages is substantially lower than in group 1, 0.68% in annual terms.

Combining the five emerging markets from Group 1 and the five from Group 2 into a 10-market portfolio (Group 3), the methodology described was applied once again. The correlation between the 100 exact and approximate semideviations is 0.93, and the average difference between them is 1.80% in annual terms. For portfolios of emerging markets, then, the proposed heuristic yields almost perfect correlations between exact and approximate semideviations, and the average differences between them are somewhat higher than for the asset classes previously discussed. In all cases, when the approximation errs, it does so on the side of caution, overestimating the risk of the portfolio in the magnitudes already discussed.

Panel C considers portfolios of individual stocks, in particular the 30 stocks from the Dow Jones Industrial Average. The 30 stocks were ordered alphabetically and split into three groups of 10 stocks. For each of these three groups, portfolios were formed following the same methodology described earlier: 100 random weights were generated for each of the 10 stocks, which were subsequently standardized to ensure that for each portfolio their sum added to one; returns for each portfolio were generated over the January 1997-December 2006 period; and their exact and approximate semideviations were calculated and subsequently annualized. As the table shows, the correlations between the exact and approximate semideviations are still very high in all three groups (0.90, 0.99, and 0.95), and the average differences between these magnitudes are 2.07%, 2.45%, and 2.34% (in annual terms) for Groups 1, 2, and 3, again higher than for asset classes.

Portfolios of the 30 Dow stocks altogether, calculated with the same methodology already described, show similar results.

The correlation between exact and approximate semideviations across the 100 portfolios remains very high (0.91); and the average difference between these two magnitudes is 2.67% in annual terms, again higher than for asset classes. As before, when the approximation errs it does so on the side of caution, overestimating the risk of the portfolio in the magnitudes already discussed.

Finally, Panel D considers portfolios of five asset classes, namely: 1) US stocks, 2) international (EAFE) stocks, 3) emerging markets stocks, 4) US bonds, and 5) US real estate. Portfolios of these five asset classes were formed with the methodology already described. The correlation between the 100 exact and approximate semideviations is a perfect 1.00; the average difference between these magnitudes, in turn, is a low 0.87% in annual terms. Again, when the approximation errs it does so on the side of caution; and again, the heuristic shows very encouraging results for asset classes.

In short, the evidence for a wide range of portfolios shows that the heuristic proposed in this article yields portfolio semivariances that are very highly correlated, as well as close in value, to the exact portfolio semivariances they intend to approximate. Importantly, as argued by Nawrocki (1999), and as is also well known, portfolio optimization is nowadays used much more for allocating funds across asset classes than across individual stocks. It is in the former case, precisely, where the heuristic approach proposed in this article is particularly accurate.

B. Optimal Portfolios

Having shown that the definition of portfolio semivariance proposed here is both simple and accurate, we can finally use it to compare the optimal portfolios that stem from mean-variance and mean-semivariance optimizations, the latter based on the proposed heuristic approach. Given that optimizers are largely used to allocate funds across asset classes, the assets considered in the optimizations are the five asset classes in Panel D of Table IV; that is, US stocks, international (EAFE) stocks, emerging markets, US bonds, and US real estate.

There are several portfolio-optimization problems, and deciding which one is more relevant simply depends on the goals and restrictions of different investors. Some may want to minimize risk; others may want to minimize risk subject to a target return; others may want to maximize return subject to a target level of risk; and others may want to maximize risk-adjusted returns. The focus of this section is on this last problem.

More precisely, the two problems considered are

$$\text{Max}_{x_1, x_2, \ldots, x_n} \quad \frac{E_p - R_f}{\sigma_p} = \frac{\sum_{i=1}^{n} x_i E_i - R_f}{\left(\sum_{i=1}^{n} \sum_{j=1}^{n} x_i x_j \sigma_{ij} \right)^{1/2}} \quad (15)$$

$$\sum_{i=1}^{n} x_i = 1 \quad \text{and} \quad x_i \geq 0 \ , \quad (16)$$

for the optimization of mean-variance (MV) portfolios, and

$$\text{Max}_{x_1, x_2, \ldots, x_n} \quad \frac{E_p - R_f}{\Sigma_{p0}} = \frac{\sum_{i=1}^{n} x_i E_i - R_f}{\left(\sum_{i=1}^{n} \sum_{j=1}^{n} x_i x_j \Sigma_{ij0} \right)^{1/2}} \quad (17)$$

$$\sum_{i=1}^{n} x_i = 1 \quad \text{and} \quad x_i \geq 0 \ , \quad (18)$$

for the optimization of mean-semivariance (MS) portfolios, where Σ_{p0} is defined as in Equation (12), and the benchmark return for the semideviation is, as before, 0%.

It is important to notice that, with the heuristic proposed here, the problem in Equations (17) and (18) can be solved with the same techniques widely used to solve the problem in Equations (15) and (16); these include professional optimization packages, simple optimization packages available with investment textbooks, and even Excel's solver. It is also important to notice that in terms of the required inputs, the only difference between these two problems is that Equations (15) and (16) require a (symmetric and exogenous) covariance matrix and Equation (17) and (18) require a (symmetric and exogenous) semicovariance matrix, which can be calculated using Equation (5).

Expected returns, required as inputs in both optimization problems, were estimated with the (arithmetic) mean return of each asset class over the whole January 1988-December 2006 sample period. Variances, covariances, semivariances, and semicovariances were calculated over the same sample period, the last two with respect to a 0% benchmark return and according to Equation (5). Optimizations were performed for combinations of three, four, and five asset classes. The results of all estimations are shown in Table V.

When optimizing a three-asset portfolio consisting of US stocks, international stocks, and emerging markets, neither the MV optimizer nor the MS optimizer give a positive weight to international stocks. Perhaps unsurprisingly, the MS optimizer gives a lower weight to emerging markets and a higher weight to the US market than does the MV optimizer. The expected monthly return of the optimal MV and MS portfolios is similar, 1.17% and 1.13%. Although the risk-adjusted return of the MS optimal portfolio is higher than that of the MV optimal portfolio, it would be deceiving to conclude that the MS optimizer outperforms the MV

Table V. Optimal Portfolios – Asset Classes

This table shows mean-variance (MV) and mean-semivariance (MS) optimal portfolios. Risk is defined as the standard deviation in MV optimizations and as the semideviation in MS optimizations. RAR denotes risk-adjusted returns defined as (Return–R_f)/Risk, where R_f denotes the risk-free rate. Return and risk are expressed in monthly terms. Monthly R_f is assumed at 0.41%. All returns are monthly, over the January 1988-December 2006 period, in dollars, and account for capital gains and dividends. All numbers but RAR are in percentages. A full data description is available in the Appendix.

	Weights					Performance		
	USA	EAFE	EMI	Bonds	NAREIT	Return	Risk	RAR
				Panel A. Three Assets				
MV	67.0	0.0	33.0			1.17	4.30	0.18
MS	78.4	0.0	21.6			1.13	2.66	0.27
				Panel B. Four Assets				
MV	17.5	0.0	18.8	63.7		0.87	2.11	0.22
MS	26.4	0.0	10.2	63.4		0.84	1.38	0.31
				Panel C. Five Assets				
MV	10.4	0.0	13.2	49.1	27.2	0.91	2.07	0.24
MS	17.1	0.0	3.0	41.4	38.5	0.90	1.41	0.35

optimizer; this is simply a consequence of the fact that the semideviation is a one-sided measure of risk.[3]

The four-asset optimization involves the previous three asset classes plus US bonds. Again both optimizers assign a zero weight to international stocks, and again the MS optimizer allocates less to emerging markets and more to the US market than does the MV optimizer. Interestingly, both optimizers allocate a substantial proportion (nearly two thirds) of the portfolio to bonds. As was the case with three assets, the expected monthly return of both optimal portfolios is very similar, 0.87% and 0.84%.

Finally, the five-asset optimization involves the previous four asset classes plus US real estate. Once again both optimizers give a zero weight to international stocks, and once again the MS optimizer allocates less to emerging markets and more to the US market than does the MV optimizer. Both optimizers allocate more than 40% of the portfolio to bonds and no less than 25% of the portfolio to real estate. And once again, the expected monthly return of both portfolios is very similar, 0.91% and 0.90%.

It is tempting to draw a conclusion regarding which optimizer performs better, but it is also largely meaningless. By definition, the MV optimizer will maximize the excess returns per unit of volatility, whereas the MS optimizer will maximize excess returns per unit of volatility below the chosen benchmark. In the end, it all comes down to what any given investor perceives as the more appropriate measure of risk.

III. An Assessment

There is little question that mean-variance optimization is far more pervasive than mean-semivariance optimization. This is, at least in part, due to the fact that mean-variance problems have well-defined, well-known closed-form solutions, which implies that users know what the optimization package is doing and what characteristics the solution obtained has. When optimizing portfolios on the basis of means and semivariances, in turn, little is usually known about the algorithms used to obtain optimal portfolios and the characteristics of the solution obtained.

This article proposes a heuristic approach for the calculation of portfolio semivariance, which essentially puts mean-semivariance optimization within reach of any academic or practitioner familiar with mean-variance optimization. By replacing the symmetric and exogenous covariance matrix by a symmetric and exogenous semicovariance matrix, the well-defined, well-known closed-form solutions of mean-variance problems can be applied to mean-semivariance problems. This takes mean-semivariance optimization away from the realm of black boxes and into the realm of standard portfolio theory.

The heuristic proposed is both simple and accurate. Estimating semicovariances is just as easy as esitmating covariances, and aggregating them into a portfolio semivariance is, with the proposed heuristic, just as easy as

[3] Comparing the risk-adjusted returns of MV and MS optimizers is non-informative at best and deceiving at worst. Similarly, comparing the efficient sets generated by these two approaches, as done by Harlow (1991) and others, yields little insight, if any. By definition, MV efficient sets will "outperform" MS efficient sets when plotted on a mean-variance graph, and the opposite will be the case when plotted on a mean-semivariance graph.

aggregating covariances into a portfolio variance. Similarly, finding optimal portfolios (regardless of whether that means minimizing risk, minimizing risk subject to a target return, maximizing return subject to a target level of risk, or maximizing risk-adjusted returns) when risk is thought of as semivariance can be done with the same methods used as when risk is thought of as variance.

In terms of accuracy, the proposed definition of portfolio semivariance was evaluated using portfolios of stocks, markets, and asset classes. The evidence discussed shows that the portfolio semivariances generated by the heuristic proposed are very highly correlated, as well as close in value, to the exact portfolio semivariances they aim to approximate. This heuristic is particularly accurate when optimizing across asset classes, which nowadays is the main use given to optimizers.

There is a growing literature on downside risk and an increasing acceptance of this idea among both academics and practitioners. Semivariance is a more plausible measure of risk than variance, as Markowitz (1991) himself suggested, and the heuristic proposed here makes mean-semivariance optimization just as easy to implement as mean-variance optimization. For this reason, this article not only provides another tool that can be added to the financial toolbox, but also hopefully contributes toward increasing the acceptance and use of mean-semivariance optimization.■

Appendix

1. A Brief Introduction to the Semideviation

This section aims to introduce the semideviation to readers largely unaware of this concept. Readers who want to explore this issue further are referred to Estrada (2006), an article from which section of the appendix borrows heavily.

A. Shortcomings of the Standard Deviation

Consider an asset with a mean annual return of 10%, and assume that in the last two years the asset returned –5% and 25%. Because both returns deviate from the mean by the same amount (15%), they both increase the standard deviation of the asset by the same amount. But is an investor in this asset equally (un)happy in both years? Not likely, which underscores one of the main problems of the standard deviation as a measure of risk: it treats an *x*% fluctuation above and below the mean in the same way, though investors obviously do not. Should it not, then, a proper measure of risk capture this asymmetry?

The second column of Table A1 shows the annual returns of Oracle (R) for the years 1995–2004. As the next-to-last row shows, the stock's mean annual return (μ) during this period was a healthy 41.1%. And as is obvious from a casual observation of these returns without resorting to any formal measure of risk, Oracle treated its shareholders to quite a bumpy ride.

The third column of the table shows the difference between each annual return and the mean annual return; for example, for the year 2004, –37.4% = 3.7%–41.1%. The fourth column shows the square of these numbers; for example, 0.1396 = $(-0.374)^2$. The average of these squared deviations from the mean is the variance (0.8418), and the square root of the variance is the standard deviation (91.7%).

Note that all the numbers in the fourth column are positive, which means that *every* return, regardless of its sign, contributes to increasing the standard deviation. In fact, the largest number in this fourth column (the one that contributes to increasing the standard deviation the most) is that for the year 1999 when Oracle delivered a *positive* return of almost 290%. Now, would an investor that held Oracle during the year 1999 be happy or unhappy? Would he count this performance *against* Oracle as the standard deviation as a measure of risk does?

We will get back to this below but before we do so consider another shortcoming of the standard deviation as a measure of risk: it is largely meaningless when the underlying distribution of returns is not symmetric. Skewed distributions of returns, which are far from unusual in practice, exhibit different volatility above and below the mean. In these cases, variability around the mean is at best uninformative and more likely misleading as a measure of risk.

B. The Semideviation

As Table A1 makes clear, one of the main problems of the standard deviation as a measure of risk is that, unlike investors, it treats fluctuations above and below the mean in the same way. However, tweaking the standard deviation so that it accounts only for fluctuations below the mean is not difficult.

The fifth column of Table A1 shows "conditional returns" with respect to the mean; that is, the lower of each return minus the mean return or zero. In other words, if a return is higher than the mean, the column shows a 0; if a return is lower than the mean, the column shows the difference between the two. To illustrate, in 1995 Oracle delivered a 44.0% return, which is higher than the mean return of 41.1%; therefore the fifth column shows a 0 for this year. In 2004, however, Oracle

Table A1. Oracle, 1995-2004

Year	R	$R-\mu$	$(R-\mu)^2$	$\text{Min}(R-\mu, 0)$	$\{\text{Min}(R-\mu, 0)\}^2$
1995	44.0%	3.0%	0.0009	0.0%	0.0000
1996	47.8%	6.7%	0.0045	0.0%	0.0000
1997	–19.8%	–60.9%	0.3709	–60.9%	0.3709
1998	93.3%	52.2%	0.2726	0.0%	0.0000
1999	289.8%	248.7%	6.1849	0.0%	0.0000
2000	3.7%	–37.3%	0.1394	–37.3%	0.1394
2001	–52.5%	–93.6%	0.8752	–93.6%	0.8752
2002	–21.8%	–62.9%	0.3952	–62.9%	0.3952
2003	22.5%	–18.6%	0.0345	–18.6%	0.0345
2004	3.7%	–37.4%	0.1396	–37.4%	0.1396
Average	**41.1%**		**0.8418**		**0.1955**
Square Root			**91.7%**		**44.2%**

delivered a 3.7% return, which is below the mean return of 41.1%; therefore, the fifth column shows the shortfall of –37.4% for this year.

Comparing the third and the fifth columns, it is clear that when a return is lower than the mean both columns show the same number; when a return is higher than the mean, however, the third column shows the difference between these two numbers and the fifth column shows a 0. Furthermore, it is clear that "conditional returns" are either negative or 0 but never positive.

The last column of Table A1 shows the square of the numbers in the fifth column. As the next-to-last row shows, the average of these numbers is 0.1955; and as the last row shows, the square root of this number is 44.2%. This number, which measures volatility *below* the mean return, is obviously a step in the right direction because we have isolated the downside that investors associate with risk. But is there anything special about the mean return? Is it possible that some investors are interested to assess volatility below the risk-free rate? Or volatility below 0? Or, more generally, volatility below *any* given return they may consider relevant?

That is exactly what the downside standard deviation of returns with respect to a benchmark B measures. This magnitude, usually referred to as the *semideviation with respect to B* (Σ_B), for short, is formally defined as

$$\Sigma_B = \sqrt{(1/T)\cdot\sum_{t=1}^{T}\{\text{Min}(R_t - B, 0)\}^2}$$

and measures *downside* volatility; or, more precisely, volatility *below* the benchmark return B. In this expression, t indexes time and T denotes the number of observations.

Table A2 shows again the returns of Oracle for the 1995-2004 period, as well as the "conditional returns" with respect to three different benchmarks: the mean return, a risk-free rate (R_f) of 5%, and 0. The third column of this exhibit is the same as the last column of Table A1 and therefore the benchmark is the mean return; the fourth and fifth columns show "conditional returns" with respect to the other two benchmarks (a risk-free rate of 5% and 0). The last row shows the semideviations with respect to all three benchmarks. (The next-to-last row shows the semivariances with respect to all three benchmarks, which are simply the square of the semideviations.)

How should these numbers be interpreted? Each semideviation measures volatility below its respective benchmark. Note that because the risk-free rate of 5% is below Oracle's mean return of 41.1%, we would expect (and find) less volatility below the risk-free rate than below the mean. Similarly, we would expect (and again find) less volatility below 0 than below the mean or the risk-free rate.

It may seem that a volatility of 21.5% below a risk-free rate of 5%, or a volatility of 19% below 0, do not convey a great deal of information about Oracle's risk. In fact, the semideviation of an asset is best used in two contexts: one is in relation to the standard deviation of the same asset and the other is in relation to the semideviation of other assets.

Table A3 shows the standard deviation (σ) of Oracle and Microsoft over the 1995-2004 period, as well as the semideviations with respect to the mean of each stock (Σ_μ), with respect to a risk-free rate of 5% (Σ_f), and with respect to 0 (Σ_0) over the same period. The semideviations of Oracle are the same as those in Table A2. The mean return of Microsoft during this period was 35.5%.

Note that although the standard deviations suggest that Oracle is far riskier than Microsoft, the semideviations tell a different story. First, note that although the volatility of Oracle below its mean is less than half of its total volatility (0.442/

Table A2. Semideviations

Year	R	$\{\mathrm{Min}(R-\mu, 0)\}^2$	$\{\mathrm{Min}(R-R_f, 0)\}^2$	$\{\mathrm{Min}(R-0, 0)\}^2$
1995	44.0%	0.0000	0.0000	0.0000
1996	47.8%	0.0000	0.0000	0.0000
1997	–19.8%	0.3709	0.0617	0.0393
1998	93.3%	0.0000	0.0000	0.0000
1999	289.8%	0.0000	0.0000	0.0000
2000	3.7%	0.1394	0.0002	0.0000
2001	–52.5%	0.8752	0.3304	0.2754
2002	–21.8%	0.3952	0.0718	0.0475
2003	22.5%	0.0345	0.0000	0.0000
2004	3.7%	0.1396	0.0002	0.0000
Average	**41.1%**	**0.1955**	**0.0464**	**0.0362**
Square Root		**44.2%**	**21.5%**	**19.0%**

Table A3. Volatility and Downside Volatility

Company	σ	Σ_μ	Σ_f	Σ_0
Oracle	91.7%	44.2%	21.5%	19.0%
Microsoft	50.4%	38.1%	23.1%	21.1%

0.917 = 48.2%), the same ratio for Microsoft is over 75% (0.381/0.504 = 75.5%). In other words, given the volatility of each stock, much more of that volatility is below the mean in the case of Microsoft than in the case of Oracle. (In fact, the distribution of Microsoft's returns has a slight negative skewness, and that of Oracle a significant positive skewness.)

Of course it is still the case that the semideviation with respect to the mean of Oracle is larger than that of Microsoft; but recall that the mean return of Oracle (41.1%) is also higher than that of Microsoft (35.5%). For this reason, it is perhaps more telling to compare semideviations with respect to the same benchmark for both stocks.

Comparing the semideviations of Oracle and Microsoft with respect to the same risk-free rate of 5%, we see that Microsoft exhibits higher downside volatility (23.1% versus 21.5%). And comparing their semideviations with respect to 0, we again see that Microsoft exhibits higher downside volatility (21.1% versus 19.0%). Therefore, although the standard deviations suggest that Oracle is riskier than Microsoft, the semideviations suggest the opposite.

2. The Data

Table A4 describes the data used in Section II of the article (The Evidence). All series are monthly, in dollars, and account for capital gains and dividends. Table 3 is based on data over the January 1997-December 2006 period and Table 4 on data over the January 1988-December 2006 period.

Table A4. The Data

USA	MSCI USA
EMI	MSCI EMI
NAREIT	FTSE NAREIT - All REITs
Emerging Markets	MSCI indices
Emerging Markets	Group 1: Five emerging markets (China, Egypt, Korea, Malaysia, Venezuela) with significant positive skewness
Emerging Markets	Group 2: Five emerging markets (Chile, Hungary, Mexico, Peru, South Africa) with significant negative skewness
Emerging Markets	Group 3: Ten markets, the five from Group 1 plus the five from Group 2
DJIA Stocks	Individual stocks from the Dow Jones Industrial Average index
DJIA Stocks	Group 1: The first ten stocks from an alphabetical ordering of the Dow (3M, Alcoa, Altria, Amex, AIG, AT&T, Boeing, Caterpillar, Citigroup, Coca-Cola)
DJIA Stocks	Group 2: The second ten stocks from an alphabetical ordering of the Dow (DuPont, ExxonMobil, GE, GM, HP, HomeDepot, Honeywell, Intel, IBM, J&J)
DJIA Stocks	Group 3: The third ten stocks from an alphabetical ordering of the Dow (JPM-Chase, McDonald's, Merck, Microsoft, Pfizer, P&G, United Tech, Verizon, WalMart, WaltDisney)
DJIA Stocks	Group 4: All thirty stocks in the Dow
Asset Classes	US stocks (MSCI USA), international stocks (MSCI EAFE), emerging markets stocks (MSCI EMI), US bonds (10-year Government bonds – Global Financial Data), and US real estate (FTSE NAREIT – All REITs)

MSCI: Morgan Stanley Capital Indices; EMI: Emerging Markets Index; FTSE: Financial Times Stock Exchange; NAREIT: National Association of Real Estate Investment Trusts; DJIA: Dow Jones Industrial Average; EAFE: Europe, Australia, and the Far East.

References

Ang, J., 1975, "A Note on the E, SL Portfolio Selection Model," *Journal of Financial and Quantitative Analysis* 10 (No. 5, December), 849-857.

Ballestero, E., 2005, "Mean-Semivariance Efficient Frontier: A Downside Risk Model for Portfolio Selection," *Applied Mathematical Finance* 12 (No. 1, March), 1-15.

De Athayde, G., 2001, "Building a Mean-Downside Risk Portfolio Frontier," in F. Sortino and S. Satchell Eds., *Managing Downside Risk in Financial Markets*, Oxford, UK, Butterworth-Heinemann.

Elton, E., M. Gruber, and M. Padberg, 1976, "Simple Criteria for Optimal Portfolio Selection," *Journal of Finance* 31 (No. 5, December), 1341-1357.

Estrada, J., 2002, "Systematic Risk in Emerging Markets: The D-CAPM," *Emerging Markets Review* 3 (No. 4, December), 365-379.

Estrada, J., 2006, "Downside Risk in Practice," *Journal of Applied Corporate Finance* 18 (No. 1, Winter), 117-125.

Estrada, J., 2007, "Mean-Semivariance Behavior: Downside Risk and Capital Asset Pricing," *International Review of Economics and Finance* 16 (No. 2), 169-185.

Hogan, W. and J. Warren, 1972, "Computation of the Efficient Boundary in the E-S Portfolio Selection Model," *Journal of Financial and Quantitative Analysis* 7 (No. 4, September), 1881-1896.

Hogan, W. and J. Warren, 1974, "Toward the Development of an Equilibrium Capital-Market Model Based on Semivariance," *Journal of Financial and Quantitative Analysis* 9 (No. 1, January), 1-11.

Grootveld, H. and W. Hallerbach, 1999, "Variance vs. Downside Risk: Is There Really that Much Difference?" *European Journal of Operational Research* 114 (No. 2, April), 304-319.

Harlow, V., 1991, "Asset Allocation in a Downside Risk Framework," *Financial Analysts Journal* 47 (No. 5, September/October), 28-40.

Markowitz, H., 1952, "Portfolio Selection," *Journal of Finance* 7 (No. 1, March), 77-91.

Markowitz, H., 1959, *Portfolio Selection: Efficient Diversification of Investments*, New York, NY, John Wiley & Sons.

Markowitz, H., 1991, *Portfolio Selection: Efficient Diversification of Investments*, 2nd ed., Cambridge, MA, Basil Blackwell.

Markowitz, H., P. Todd, G. Xu, and Y. Yamane, 1993, "Computation of Mean-Semivariance Efficient Sets by the Critical Line Algorithm," *Annals of Operations Research* 45 (No. 1, December), 307-317.

Nawrocki, D., 1983, "A Comparison of Risk Measures When Used in A Simple Portfolio Selection Heuristic," *Journal of Business Finance and Accounting* 10 (No. 2, June), 183-194.

Nawrocki, D., 1999, "A Brief History of Downside Risk Measures," *Journal of Investing* (Fall), 9-25.

Nawrocki, D. and K. Staples, 1989, "A Customized LPM Risk Measure for Portfolio Analysis," *Applied Economics* 21 (No. 2, February), 205-218.

Sing, T. and S. Ong, 2000, "Asset Allocation in a Downside Risk Framework," *Journal of Real Estate Portfolio Management* 6 (No. 3), 213-223.